Columbia University in the City of New York
Department of Computer Science (212) 854-2736
New York, NY 10027 tech-reports@cs.columbia.edu

TECHNICAL REPORTS LIST

The Computer Science Department at Columbia University is pleased to announce the enclosed list of technical reports and abstracts for 1992.

If you wish to order any of these reports, please complete the form below, fold so that the return address is properly positioned, and affix postage.

Please, send us your electronic mail address to be included in our electronic distribution list. Send your electronic address to `tech-reports@cs.columbia.edu`.

A NOTE TO INTERNET USERS: many Columbia Computer Science technical reports are available by anonymous FTP from host `cs.columbia.edu` directory `pub/reports`. All reports are stored in compressed postscript format.

Technical Report Numbers

1. ____	2. ____	3. ____	4. ____	5. ____
6. ____	7. ____	8. ____	9. ____	10. ____
11. ____	12. ____	13. ____	14. ____	15. ____
16. ____	17. ____	18. ____	19. ____	20. ____
21. ____	22. ____	23. ____	24. ____	25. ____
26. ____	27. ____	28. ____	29. ____	30. ____
31. ____	32. ____	33. ____	34. ____	35. ____
36. ____	37. ____	38. ____	39. ____	40. ____
41. ____	42. ____	43. ____	44. ____	45. ____
46. ____	47. ____	48. ____	49. ____	50. ____
51. ____	52. ____	53. ____	54. ____	55. ____
56. ____	57. ____	58. ____		

Name ________________________________

Address ________________________________

City ____________ State ____________ Zip ______

Attn:Technical Reports Librarian
Department of Computer Science
Columbia University
New York, NY 10027

[CUCS-001-92] Concurrency control in rule-based software development environments (ph.d. thesis).

Naser S. Barghouti.

This dissertation investigates the concurrency control problem in software development environments (SDEs). The problem arises when multiple developers perform activities that concurrently access the project's components, stored as database objects. The interleaved execution of the developers' activities leads to interference if they access overlapping sets of objects concurrently. An SDE can ensure that activities never interfere by modeling their execution in terms of atomic transactions and allowing only serializable schedules. This prevents cooperation, which requires some degree of interference between the activities of multiple developers. To allow cooperation, an SDE must be provided with semantic information about development activities. In rule-based SDEs, the necessary information is readily available in the set of rules that defines the process model of a project. The rules are loaded into the SDE, which provides process-specific assistance through a rule chaining engine. A single user command might lead the chaining engine to initiate a rule chain. The concurrency control problem in rule-based SDEs manifests itself in terms of interference between concurrent rule chains. We present a mechanism that extracts semantic information from the process model to solve the concurrency control problem without obstructing cooperation. The mechanism is composed of two modules: (1) a conflict detection module, which models activities as nested transactions and uses two-phase locking to detect interference; and (2) a conflict resolution module, which employs two protocols to resolve interference. The first protocol, SCCP, uses the process model to implement a priority-based scheme that aborts the "least important" of the interfering transactions. The second protocol, PCCP, overrides SCCP by consulting process-specific control rules, written by the project administration. Each control rule describes a specific interference and the actions that resolve it. We have implemented SCCP and parts of PCCP in @marvel, a multi-user rule-based SDE developed at Columbia.

[CUCS-002-92] Parallel string matching algorithms.

Dany Breslauer and Zvi Galil.

The string matching problem is one of the most studied problems in computer science. While it is very easily stated and many of the simple algorithms perform very well in practice, numerous works have been published on the subject and research is still very active. In this paper we survey recent results on parallel algorithms for the string matching problem.

[CUCS-003-92] A lower bound for parallel string matching.

Dany Breslauer and Zvi Galil.

We present an $\Omega(\log\log m)$ lower bound on the number of rounds necessary for finding occurrences of a pattern string $P[1..m]$ in a text string $T[1..2m]$ in parallel using m comparisons in each round. The bound is within a constant factor of the fastest algorithm for this problem and also holds for an m-processor CRCW-PRAM in the case of a general alphabet. Consequently we derive the parallel complexity of the string matching problem using p processors for general alphabets which is,

- $\Theta(\frac{m}{p})$ if $p \leq \frac{m}{\log\log m}$
- $\Theta(\log\log m)$ if $\frac{m}{\log\log m} \leq p \leq m$
- $\Theta(\log\log_{\frac{2p}{m}} p)$ if $m \leq p \leq m^2$
- $\Theta(1)$ if $p \geq m^2$

or in short $\Theta(\lceil\frac{m}{p}\rceil + \log\log_{\lceil 1+p/m\rceil} 2p)$.

[CUCS-004-92] A survey of image registration techniques.

Lisa Gottesfeld Brown.

Registration is a fundamental task in image processing used to match two or more pictures taken, for example, at different times, from different sensors or from different viewpoints. Over the years, a broad range of techniques have been developed for the various types of data and problems. These techniques have been independently studied for several different applications resulting in a large body of research. This paper organizes this material by establishing the relationship between the distortions in the image and the type of registration techniques which are most suitable. Two major types of distortions are distinguished. The first type are those which are the source of misregistration, i.e., they are the cause of the misalignment between the two images. Distortions which are the source of misregistration determine the transformation class which will optimally align the two images. The transformation class in turn influences the general technique that should be taken. The second type of distortion are those which are not the source of misregistration; this type usually effects intensity values but they may also be spatial. Distortions of this type are not to be removed by registration but they make registration more difficult since an exact match is no longer possible. They effect the choice of feature space, similarity measure and search space and strategy which will make up the final technique. All registration techniques can be viewed as different combinations of these choices. This framework is useful for understanding the merits and relationships between the wide variety of existing techniques and for assisting in the selection of the appropriate technique for a specific problem.

[CUCS-005-92] Extraction and use of contextual attributes for theory completion: An integration of explanation-based and similarity-based learning.

Andrea Pohoreckyj Danyluk.

This research investigates the use of contextual cues to address problems in machine learning that arise from assumptions about the initial knowledge that is necessary for the acquisition of new information. Machine learning approaches may be placed along a spectrum describing purely inductive to purely deductive techniques. Inductive systems possess essentially no explicit knowledge that can be used in acquiring new facts, while deductive systems are assumed to contain a complete theory of the domain. Most work in machine learning has concentrated on approaches at the two ends of the spectrum. This dissertation describes an approach that integrates inductive and deductive methods. It provides a mechanism by which induction can be used in order to detect and acquire knowledge missing from the domain theory of a deductive system, while extracting information to guide the inductive search for missing knowledge. This information may be termed @i(contextual). Much knowledge can be brought to bear in filling gaps in a deductive system's domain theory. This might include information about the domain theory itself or perhaps the history of the learning system. I show that a domain-independent set of attributes can be used to describe contextual information available in the deductions made by a system. Furthermore, I present an algorithm for using that information to generate domain-independent heuristics that guide induction. The heuristics use context information to select examples for induction of missing information and guide the way the examples are interpreted. My empirical investigations have yielded a characterization of the efficacy of various subsets of attributes. The algorithm for extracting and applying contextual knowledge is implemented in the Gemini system. Gemini combines similarity-based learning, an inductive technique, and explanation-based learning, a deductive method, in an architecture of mutual dependence. Similarity-based learning is used to fill gaps detected in the domain theory used by explanation-based learning. I have developed Gemini as a general learning system and have tested it in four domains. It has been the testbed for my investigation of contextual attributes and confirms that there exists a set of domain-independent attributes that can effectively guide an inductive search for concept definitions.

[CUCS-006-92] A structural classification of integrated replica control mechanisms.

Shu-Wie Chen and Calton Pu.

We classify practical replica control (RC) mechanisms according to two main criteria: (1) the abstract RC protocol, e.g., how many physical copies are written for a logical update, and (2) the kind of integration between the RC mechanism and system services such as atomic broadcast, commit protocols, and concurrency control. We describe about a dozen representative RC mechanisms (most of them implemented) classified by their integration with communications, concurrency control, or both. This classification based on a few essential components of RC mechanisms show that practical replicated systems overwhelmingly adopt the read-one/write-one RC protocol (read/write one physical copy for each logical read/write), due to its low overhead and high availability, despite some inherent inconsistency. Furthermore, this trend points to the need for the integration of weak consistency methods (e.g., epsilon serializability) with read-one/write-one RC protocols if replicated databases are to be realized in practice.

[CUCS-007-92] An efficient variable-consistency replicated file service.

Carl D. Tait and Dan Duchamp.

In a previous paper, we introduced a new method of file replica management designed to address the problems faced by mobile clients. Two goals shaped our design: minimizing synchronous operations that involve multiple servers, and letting clients determine the level of consistency desired. The previous paper was a high-level view of the most fundamental issues; here, we refine our previous ideas and flesh out the rest of the design.

[CUCS-008-92] Secure distributed computing.

Matthew K. Franklin.

Secure distributed computing protocols allow a group of players, within some specific computational environment, to evaluate jointly the output of a function while maintaining the secrecy of privately held inputs. An oft-cited example is secret ballot election, in which a final tally must be computed without revealing anything else about the individual votes. The secrecy constraints of more general protocol problems can be satisfied by incorporating secure distributed computing (either cryptographic or unconditional) as a basic building block. This paper surveys the extensive progress that has been made in the area of secure distributed computation over the past few years.

[CUCS-009-92] Communication complexity of secure computation.

Matthew K. Franklin and Moti Yung.

A secret-ballot vote for a single proposition is an example of a secure distributed computation. The goal is for n participants to jointly compute the output of some n-ary function (in this case, the sum of the votes), while protecting their individual inputs against some form of misbehavior. In this paper, we initiate the investigation of the communication complexity of unconditionally secure multi-party computation, and its relation with various fault-tolerance models. We present upper and lower bounds on communication, as well as tradeoffs among resources. First, we consider the "direct sum problem" for communication complexity of perfectly secure protocols: Can the communication complexity of securely computing a single function $f : F^n \rightarrow F$ at k sets of inputs be smaller if all are computed simultaneously than if each is computed individually? We show that the answer depends on the failure model. A factor of $O(\frac{n}{\log n})$ can be gained in the privacy model (where processors are curious but correct); specifically, when f is n-ary addition (mod 2), we show a lower bound of $\Omega(n^2)$ bits for computing f once and an upper bound of $O(n^2 \log n)$ for computing f $O(n)$ times simultaneously. No gain is possible in a slightly stronger fault model (fail-stop mode); specifically, when f is n-ary addition over $GF(q)$, we show an exact bound of $\Theta(kn^2 \log q)$ for computing f at k sets of inputs simultaneously (for any $k \geq 1$). However, if one is willing to pay an additive cost in fault tolerance (from t to $t-k+1$), then a variety of known non-cryptographic protocols (including "provably unparallelizable" protocols from above!) can be systematically compiled to compute one function at k sets of inputs *with no increase in communication complexity.* Our compilation technique is based on a new compression idea of polynomial-based multi-secret sharing. Lastly, we show how to compile private protocols into error-detecting protocols at a big savings of a factor of $O(n^3)$ (up to a log factor) over the best known error-correcting protocols. This is a new notion of fault-tolerant protocols, and is especially useful when malicious behavior is infrequent, since error-detection implies error-correction in this case.

[CUCS-010-92] Meldc threads: Supporting large-scale dynamic parallelism.

Wenwey Hseush, James C. Lee, and Gail Kaiser.

We present a thread model that supports large-scale dynamic parallelism for a concurrent object-oriented language, @meldc. Large-scale dynamic parallelism indicates that the number of threads that co-exist is large and the life-time for each thread is short. We introduce the concept of interleaving stack (IS), which melds frames of multiple threads into one shared stack. The pure interleaving stack is not scalable due to severe internal fragmentation. We design and evaluate two versions, circular interleaving stack (CIS) and @i(k)-interleaving stack (@i(k)-IS). Our simulation shows that in the domains of large-scale dynamic parallelism, CIS performs better than the traditional one-stack-per-thread (1SPT) mechanism with respect to memory utilization. The 1SPT mechanism gains better memory utilization when thread life-time increases. In the domains where most threads are short lived and few are long lived, @i(k)-IS shows better memory utilization than others. CIS, @i(k)-IS and 1SPT show similar CPU utilization. We also introduce the notion of thread mobility. Our analysis shows less than 3threads have better chance to survive stack overflow. Once a fault of stack overflow occurs to a thread, it can move to a bigger memory space or garbage collect memory fragments and resume its computation. The IS mechanisms have higher survivability than 1SPT, since they allow individual frames to be relocated, rather than moving the entire stack.

[CUCS-011-92] Computers without clocks.

Stephen H. Unger.

Continuing increases in logic density on VLSI chips has led to increasing problems with clock distribution systems. This has stimulated interest in dispensing with the clocks that control virtually all existing digital systems. The recent construction at Caltech of a microprocessor chip without a clock has clearly demonstrated the feasibility of such systems. Basic principles that can be used to design computers without clocks are outlined here. Handshaking and dual-rail coding constitute one important related pair of concepts. Logic circuit level designs of asynchronous registers, counters, shift registers, and adders are shown. Control modules are described that can be used as building blocks for systems of varying complexity. These can be used to implement all of the basic features of modern computers including interrupts and pipelines. The design of an asynchronous add-and-shift binary multiplier is used to illustrate the use of these data processing and control modules. Most of the work shown uses what is generally referred to as 4-phase handshaking, but 2-phase handshaking is also discussed. The extent to which unclocked systems can be truly delay-insensitive is discussed.

[CUCS-013-92] Rule-based process servers for software development environments.

Steve Popovich.

Current software development environments generally take one of two approaches. The integrated toolset approach provides an extensible set of utilities to the developer, and a common inter-tool communication and integration facility. However, it provides no guidance as to how the tools should be used. The process-based approach integrates a toolset with a process control engine so as to provide an environment that enacts a particular development process. Unfortunately, it is not always feasible to go the process-based route; in the case where the existing tools are integrated with one another, the process-based environment's integration mechanism may be in conflict with that of the integrated toolset. Process servers are an alternative means of integrating a process control engine with a toolset, which avoids this conflict by adapting the process control engine to use the toolset's existing integration mechanism. Many process-based environments are rule-based environments: Their process specification language is a rule language. Rules deal well with local constraints, but often interact counter-intuitively and make it difficult to express the global nature of a process. Higher-level process languages are needed to describe global constraints effectively but without losing the advantages of rules. We propose to integrate the two approaches in an environment centered around a @i(rule-based process server), which is a process server with a rule-based control engine. We propose also to provide a framework for implementing declarative process modeling languages and for interfacing the interpreter to the target toolset. Although our process server is to be rule-based, we plan to support a broader, multi-paradigm notion of a process modeling language. These process modeling languages will include declarative global process specifications in addition to local rule-based constraints, but translated into a rule-based form, plus appropriate run-time support wrapping around the rule interpreter, for execution.

[CUCS-014-92] Resource freshness in real-time scheduling.

Shyhtsun F. Wu and Gail E. Kaiser.

We define and study a *real-time producers/consumers* problem, where the production/consumption relations are *nondeterministic*. In other words, a consumer can get one unit of resources from any one of the producers, and one producer can provide one unit of resources for any one of the consumers. Furthermore, a producer for one type of resources might be the consumer for another type of resources. A class of *fresh resource allocation policies* is developed not only to make resources *available* for periodic real-time tasks but also to control the *resource freshness*. We first model the real-time scheduling problem as a feasible *minimum cost flow (MCF)* problem. Then, by defining the *linear cost function* differently, we get resource allocation policies for making trade-off between *freshness* and *schedulability*.

Finally, we show that our result can be used to extend the *priority ceiling protocol (PCP)* to work with *general unbounded semaphores*, where *Wait* and *Signal* operations need not be in pairs in one task.

[CUCS-015-92] Navigation by isolated landmarks.

Il-Pyung Park.

In this paper we propose and demonstrate a simple method for navigation in a large unstructured environment that contains featureless objects, which uses "isolated" landmarks in the navigator's view. The map-maker and navigator are implemented using an IBM 7575 SCARA robot arm, the PIPE (Pipelined Image Processing Engine), and two cameras. The navigational environment consists of a flat plane with spherical objects populated randomly on it. First, the map-maker model observes the environment, given a starting position and a goal position, and it generates a "custom map" that describes describes "how" to get from the starting position to the goal position. The accuracy and the efficiency of the directional instructions are then demonstrated by the navigator by following the commands in the "custom map". This is a first step towards our eventual goal, which is to develop a full set of vocabulary that can qualitatively "describe" the navigational environment to direct the navigator with efficiency and accuracy.

[CUCS-016-92] A bi-level language for software process modeling.

Gail E. Kaiser, Steven S. Popovich, and Israel Z. Ben-Shaul.

Most software process modeling formalisms are well-suited to one of two levels of specification, but not both. Some process languages concentrate on @i(global) control flow and synchronization; these languages make it easy to define the broad outline of a process, but harder (or even impossible) to refine the process by expressing constraints and policies on individual tools and data. Other process formalisms are inherently @i(local): it is easy to define constraints, but far from straightforward to express control flow. We propose combining global and local formalisms to produce single but bi-level process modeling languages. We present our bi-level Activity Structures Language, which combines constrained expressions with rules, and its implementation on top of the @marvel 3.0 multi-user environment.

[CUCS-017-92] Finding all periods and initial palindromes of a string in parallel.

Dany Breslauer and Zvi Galil.

An optimal $O(\log\log n)$ time CRCW-PRAM algorithm for computing all periods of a string is presented. Previous parallel algorithms compute the period only if it is shorter than half of the length of the string. This algorithm can be used to find all initial palindromes of a string in the same time and processor bounds. Both algorithms are the fastest possible over a general alphabet. We derive a lower bound for finding palindromes by a modification of a previously known lower bound for finding the period of a string. When p processors are available the bounds become $\Theta(\lceil\frac{n}{p}\rceil + \log\log_{\lceil 1+p/n\rceil} 2p)$.

[CUCS-018-92] Towards provably secure efficient electronic cash.

Matthew Franklin and Moti Yung.

[CUCS-019-92] Kuratowski-type theorems for average genus.

Jainer Chen and Jonathan L. Gross.

Graphs of small average genus are characterized. In particular, a Kuratowski-type theorem is obtained: except for finitely many graphs, a cutedge-free graph has average genus less than or equal to 1 if and only if it is a necklace. We provide a complete list of those exceptions. A Kuratowski-type theorem for graphs of maximum genus 1 is also given. Some of the methods used in obtaining these results are extensions of a classical result of Whitney.

[CUCS-020-92] Limit points for average genus (i): 3-connected and 2-connected simplical graphs.

Jainer Chan and Jonathan L. Gross.

It is demonstrated that a given value of average genus is shared by at most finitely many 2-connected simplicial graphs and by at most finitely many 3-connected graphs. Moreover, the distribution of values of average genus is sparse, in the following sense: within any finite real interval, there are at most finitely many different numbers that are values of average genus for 2-connected simplicial graphs or for 3-connected graphs. Thus, there are no limit points for the values of average genus of graphs in these classes. The potential applicability of these results to graph isomorphism testing is conisdered.

[CUCS-021-92] Limit points for average genus (ii): 2-connected non-simplicial graphs.

Jainer Chen and Jonathan L. Gross.

In part (I) of this paper, it is proved that there are no limit points for the set of values of average genus of 2-connected simplicial graphs and of 3-connected graphs. The need for such restrictions is now demonstrated by showing that infinitely many real numbers are limit points of values of average genus for 2-connected non-simplicial graphs. A systematic method for constructing such limit points is presented, and it is proved that this method is essentially the only way to construct limit points of values of average genus for "homeomorphically nested" 2-connected graphs.

[CUCS-022-92] An on-line string superprimitivity test.

Dany Breslauer.

A string w *covers* another string z if every symbol of z is within some occurrence of w in z. A string is called *superprimitive* if it is covered only by itself, and *quasiperiodic* if it is covered by some shorter string. We present an on-line linear-time algorithm that tests if each prefix of an input string is superprimitive while the string is given a symbol at a time.

[CUCS-023-92] Extending user expertise in interactive environments: A task-centered approach to automatic assistance.

Ursula Caroline Wolz.

Interactive environments present users with inherent conflicts between ease of use and functional power and extensibility. Learning to exploit that power can be difficult. A critical problem for Computer Science is how users can effectively extend their expertise without suspending the task at hand. This thesis addresses the problem of providing personalized assistance to a user through an on-line question answering system called G@c[enie]. G@c[enie] generates textual information calibrated both to the task at hand and the user's past experience within the environment. The work encompasses natural language generation, knowledge representation and planning for task-centered settings, and intelligent computer assisted instruction. This work is distinctive in that G@c[enie] explicitly follows both linguistic and pedagogical principles, first by responding informatively and second by opportunistically attempting to enrich users' knowledge. Three kinds of expertise that exploit a three part user model are necessary to accomplish this. The expertise is characterized as domain expertise, analytic expertise, and explanatory expertise. The user model consists of a situational context, a discourse context and a model of user task expertise. This work extends natural language generation by unifying the structural advantage of ATN-based schemata with the flexibility of classic text planning. It extends knowledge representation and planning in task-centered settings by clearly separating knowledge of @i(intentions), that is how domain specific goals may be accomplished, and @i(causality), that is, what the results of actions are. Furthermore it provides a mechanism by which semantic distinctions between plans can be encoded. Finally it presents a unique approach to intelligent computer assisted instruction since it does not encode a curriculum, but bases an instructional episode on the task at hand.

[CUCS-024-92] Efficient string algorithmics (ph.d. thesis).

Dany Breslauer.

Problems involving strings arise in many areas of computer science and have numerous practical applications. We consider several problems from a theoretical perspective and provide efficient algorithms and lower bounds for these problems in sequential and parallel models of computation. In the sequential setting, we present new algorithms for the string matching problem improving the previous bounds on the number of comparisons performed by such algorithms. In parallel computation, we present tight algorithms and lower bounds for the string matching problem, for finding the periods of a string, for detecting squares and for finding initial palindromes.

[CUCS-025-92] A survey of empirical measurements of networked software systems.

Frederick M. Korz.

Empirical measurement is an important research tool in networked systems, providing concrete answers to questions including latency, bandwidth, utilization, and communications patterns. The presence of networks, however, complicates the task of measuring software systems. Networks introduce a host of often uncontrollable factors – message latency, network reliability, etc. We examine issues in empirical measurement, particularly in the measurement of networked software systems. We then present a survey of representative measurement literature using a framework derived from these measurement issues. Finally we suggest the probable direction of future work in the field.

[CUCS-026-92] Estimating the largest eigenvalue by the power and lanczos algorithms with a random start.

J. Kuczynski and H. Wozniakowski.

Our problem is to compute an approximation to the largest eigenvalue of an $n \times n$ large symmetric positive definite matrix with relative error at most ε. We consider only algorithms that use Krylov information $[b, Ab, \ldots, A^k b]$ consisting of k matrix-vector multiplications for some unit vector b. If the vector b is chosen deterministically then the problem cannot be solved no matter how many matrix-vector multiplications are performed and what algorithm is used. If, however, the vector b is chosen randomly with respect to the uniform distribution over the unit sphere, then the problem can be solved on the average and probabilistically. More precisely, for a randomly chosen vector b we study the power and Lanczos algorithms. For the power algorithm (method) we prove sharp bounds on the average relative error and on the probabilistic relative failure. For the Lanczos algorithm we present only upper bounds. In particular, $\ln(n)/k$ characterizes the average relative error of the power algorithm, whereas $O((\ln(n)/k)^2)$ is an upper bound on the average relative error of the Lanczos algorithm. In the probabilistic case, the algorithm is characterized by its probabilistic relative failure which is defined as the measure of the set of vectors b for which the algorithm fails. We show that the probabilistic relative failure goes to zero roughly as $\sqrt{n}(1-\varepsilon)^k$ for the power algorithm and at most as $\sqrt{n}\, e^{-(2k-1)\sqrt{\varepsilon}}$ for the Lanczos algorithm. These bounds are for a worst case distribution of eigenvalues which may depend on k. We also study the behavior in the average and probabilistic cases of the two algorithms for a fixed matrix A as the number of matrix-vector multiplications k increases. The bounds for the power algorithm depend then on the ratio of the two largest eigenvalues and their multiplicities. The bounds for the Lanczos algorithm depend on the ratio between the difference of the two largest eigenvalues and the difference of the largest and the smallest eigenvalues.

[CUCS-027-92] Probabilistic bounds on the extremal eigenvalues and condition number by the lanczos algorithmprobabilistic bounds on the extremal eigenvalues and condition number by the lanczos algorithm.

J. Kuczynski and H. Wozniakowski.

We analyze the Lanczos algorithm with a random start for approximating the extremal eigenvalues of a symmetric positive definite matrix. We present some bounds on the Lebesgue measure (probability) of the sets of these starting vectors for which the Lanczos algorithm gives at the k-th step satisfactory approximations to the largest and smallest eigenvalues. Combining these bounds we get similar estimates for the condition number of a matrix.

[CUCS-028-92] Relaxed verification for continuous problems", to appear in journal of complexity.

E. Novak and H. Wozniakowski.

We analyze the complexity of verifying whether a given element is close to a solution element. Closeness is measured by two nonnegative parameters, ε and α. If $\alpha = 0$, then we get the strong verification problem which usually cannot be solved in the worst case setting, see Wozniakowski [90]. For $\alpha > 0$, we have the relaxed verification problem which is studied in this paper in the worst case setting. We show that the (ε, α)-verification complexity is roughly equal to the η-computation complexity, where $\eta = \eta(\varepsilon, \alpha)$. Under certain general assumptions we prove that $\eta(\varepsilon, \alpha) = \Theta(\varepsilon\, \alpha^r)$, where r depends on the problem norm, $r \in [0, 1]$. We also show that adaption sometimes helps significantly.

[CUCS-029-92] Measures of uncertainty and information in computation", to appear in journal of information science.

E. W. Packel, J. F. Traub, and H. Wozniakowski.

[CUCS-030-92] The monte carlo algorithm with a pseudo-random generator.

J. F. Traub and H. Wozniakowski.

We analyze the Monte Carlo algorithm for the approximation of multivariate integrals when a pseudo-random generator is used. We establish lower and upper bounds on the error of such algorithms. We prove that as long as a pseudo-random generator is capable of producing only finitely many points, the Monte Carlo algorithm with such a pseudo-random generator fails for L_2 or continuous functions. It also fails for Lipschitz functions if the number of points does not depend on the number of variables. This is the case if a linear congruential generator is used with one initial seed. On the other hand, if a linear congruential generator of period m is used for each component with independent uniformly distributed initial seeds, then the Monte Carlo algorithm with such a pseudo-random generator using n function values behaves as for the uniform distribution and its expected error is roughly $n^{-1/2}$ as long as the number n of function values is less than m^2.

[CUCS-031-92] Complexity of verification and computation for ibc problems", to appear in journal of complexity.

H. Wozniakowski.

We analyze the complexity of *verifying* whether a given element is within ε of the solution element. This may be contrasted with the complexity of *computing* an element that is within ε of the solution element. For discrete problems with $\varepsilon = 0$ verification is no harder than computation in any setting. For IBC problems verification can be easier or harder than computation. We will show that the worst case complexity of verification for IBC problems is often infinite. We therefore switch to the probabilistic case and study the probabilistic complexity of verification as a function of the error tolerance ε and the probability of failure δ. We assume that the solution element is specified by a linear continuous functional defined on a Banach space equipped with a Gaussian measure. For fixed δ and small ε, the complexity of verification is zero, whereas for fixed ε and small δ the complexity of verification is essentially a function of only δ and maybe exponentially harder than the complexity of computation.

[CUCS-032-92] Average case complexity of linear multivariate problems. part i: Theory.

H. Wozniakowski.

We study the average case complexity of linear multivariate problems, that is, the approximation of continuous linear operators on functions of d variables. The function spaces are equipped with Gaussian measures. We consider two classes of information. The first class Λ^{std} consists of function values, and the second class Λ^{all} consists of all continuous linear functionals. Tractability of a linear multivariate problem means that the average case complexity of computing an ε-approximation is $O\left((1/\varepsilon)^p\right)$ with p independent of d. The smallest such p is called the exponent of the problem. Under mild assumptions, we prove that tractability in Λ^{all} is equivalent to tractability in Λ^{std}, and that the difference of the exponents is at most 2. The proof of this result is not constructive. We provide a simple condition to check tractability in Λ^{all}. We also address the issue how to construct optimal (or nearly optimal) sample points for linear multivariate problems. We use relations between average case and worst case settings. These relations reduce the study of the average case to the worst case for a different class of functions. In this way we show how optimal sample points from the worst case setting can be used in the average case. In Part II we will apply the theoretical results to obtain optimal or almost optimal sample points, optimal algorithms and average case complexity functions for linear multivariate problems equipped with the folded Wiener sheet measure. Of particular interest will be the multivariate function approximation problem.

[CUCS-033-92] Average case complexity of linear multivariate problems. part ii: Applications.

H. Wozniakowski.

We apply the theoretical results of Part I to linear multivariate problems LMP equipped with the folded Wiener sheet measure. We are particularly interested in multivariate weighted integration and multivariate function approximation. We prove that any LMP which satisfies (A.1) of Part I is tractable and its exponent is at most 2. We show that optimal or nearly optimal sample points can be derived from hyperbolic cross points, and exhibit nearly optimal algorithms. In particular, we find the order of the average case complexity of multivariate function approximation in Λ^{std}.

[CUCS-034-92] Visual learning and recognition in parametric eigenspace.

Hiroshi Murase and Shree K. Nayar.

We address the problem of automatically learning object models for object recognition and pose estimation. We approach the recognition problem as one of appearance matching rather than shape matching. The appearance of an object in a two-dimensional image depends on its shape, reflectance properties, pose in the scene, and the illumination conditions. While the first two are intrinsic properties of the object, pose and illumination vary from scene to scene. For each object of interest to the recognition system, a large set of images is obtained by varying pose and illumination in small increments. This large image set is compressed using the Karhunen-Loeve method to obtain a low-dimensional eigenspace in which the object is represented as a parametric hypersurface. We refer to this compact representation of object appearance as the parametric eigenspace. The parameters of the representation are pose and illumination direction. Given an unknown input image, the recognition system projects the image onto the parametric eigenspace. The object is recognized based on the hypersurface it lies on. The exact position of the projection on the hypersurface determines the object's pose in the image. While the learning stage requires the acquisition of a large number of object images and the computation of the eigenvectors of object image set, the recognition and pose estimation stages are simple and computationally efficient. We have conducted experiments using four sample objects. The performance of the recognition and pose estimation algorithms was analyzed using over a thousand input images of the four objects. The sensitivity of recognition to the number of eigenspace dimensions and the number of learning samples used is analyzed. We have found that object representation in eigenspaces with less than 10 dimensions produces accurate recognition results with an average pose estimation error of 0.5 degrees. These results suggest the parametric eigenspace to be a valuable learning tool for vision applications varying from object recognition to visual inspection.

[CUCS-035-92] Linear ill-posed problems are solvable on the average for all gaussian measures.

J. F. Traub and A. G. Werschulz.

[CUCS-036-92] Finiteness of the set of frames of graphs of bounded average genus.

Jianer Chen and Jonathan L. Gross.

In this report, a detailed proof is given that for every real number r, there is a finite set $\mathcal{F}_{<r}$ of graphs, which are called "frames", such that every cutedge-free smooth graph whose average genus is less than r can be obtained by attaching a set of non-overlapping ears to some member of $\mathcal{F}_{<r}$.

[CUCS-037-92] Average case complexity of linear multivariate problems. part ii: Applications.

H. Wozniakowski.

We apply the theoretical results of Part I to linear multivariate problems LMP equipped with the folded Wiener sheet measure. We are particularly interested in multivariate weighted integration and multivariate function approximation. We prove that any LMP which satisfies (A.1) of Part I is tractable and its exponent is at most 2. We show that optimal or nearly optimal sample points can be derived from hyperbolic cross points, and exhibit nearly optimal algorithms. In particular, we find the order of the average case complexity of multivariate function approximation in Λ^{std}.

[CUCS-038-92] Planning in ai and text planning in natural language generation.

Jong Gyun Lim.

Researchers in natural language generation are becoming increasingly aware of the similarities between the task of planning the content and structure of the natural language text and that of other AI planning tasks. The problem of text planning and other AI planning problems have been studied separately from each other, and while many overlapping results have been found, they have not been shared as often. The purpose of this paper is to explore the similarities between the two areas of research by reviewing the problems and solutions of each in order to gain a better insight into how needs in one area can benefit from results in another.

[CUCS-039-92] Synthesis: An efficient implementation of fundamental operating system services.

Henry Massalin.

This dissertation shows that operating systems can provide fundamental services an order of magnitude more efficiently than traditional implementations. It describes the implementation of a new operating system kernel, Synthesis, that achieves this level of performance. The Synthesis kernel combines several new techniques to provide high performance without sacrificing the expressive power or security of the system. The new ideas include: + Run-time code synthesis– a systematic way of creating executable machine code at runtime to optimize frequently-used kernel routines – queues, buffers, context switchers, interrupt handlers, and system call dispatchers – for specific situations, greatly reducing their execution time. + Fine-grain scheduling– a new process-scheduling technique based on the idea of feedback that performs frequent scheduling actions and policy adjustments (at sub-millisecond intervals) resulting in an adaptive, self-tuning system that can support real-time data streams. + Lock-free optimistic synchronization is shown to be a practical, efficient alternative to lock-based synchronization methods for the implementation of multiprocessor operating system kernels. + An extensible kernel design that provides for simple expansion to support new kernel services and hardware devices while allowing a tight coupling between the kernel and the applications, blurring the distinction between user and kernel services. The result is a significant performance improvement over traditional operating system implementations in addition to providing new services.

[CUCS-040-92] An optimal $o(\log\log n)$ time parallel algorithm for detecting all squares in a string.

Alberto Apostolico and Dany Breslauer.

An optimal $O(\log\log n)$ time CRCW-PRAM algorithm for detecting all squares in a string is presented. A tight lower bound shows that over a general alphabet this is the fastest possible optimal algorithm. When p processors are available the bounds becomes $\Theta(\lceil\frac{n\log n}{p}\rceil + \log\log_{\lceil 1+p/n\rceil} 2p)$.

[CUCS-041-92] Fast parallel string prefix-matching.

Dany Breslauer.

An $O(\log\log m)$ time $\frac{n\log m}{\log\log m}$-processor CRCW-PRAM algorithm for the string prefix-matching problem over a general alphabet is presented. The algorithm can also be used to compute the KMP failure function in $O(\log\log m)$ time on $\frac{m\log m}{\log\log m}$ processors. These results improve on the running time of the best previous algorithm for both problems, which was $O(\log m)$, while preserving the same number of operations.

[CUCS-042-92] Generating newswire report leads with historical information: a draft and revision approach.

Jacques Robin.

[CUCS-043-92] Tailoring lexical choice to the user's vocabulary in multimedia explanation generation.

Jacques Robin.

[CUCS-044-92] Meldc runtime kernel implementation guide.

James C. Lee, Wenwey Hseush, and Scott Meyer (Writer).

This document is an implementation manual written for the MeldC environment. The document descibes the necessary code needed to implement the MeldC environment so that it can be used by users looking for an object oriented environment. This code includes the following major sections: 1) The Kernel environment used by MeldC. 2) The meldC_user.h file used to define many parts of MeldC. 3) The main class used by MeldC, the MetaClass. 4) The Shadow and Persistent object concepts used by MeldC. 5) The Runtime code needed by MeldC to carry out many of the principles used by MeldC Each of these are discussed separately in the manual in their respective chapters. After reading this manual, we hope that you will understand the MeldC environment well enough to use it to it's limits.

[CUCS-045-92] Design & implementation of the meldc translator.

Esther Y. Woo.

This report discuss the design and implementation of the compiler for the MeldC programming language. It is intended to serve as an implementor's guide for the development and maintenance of the MeldC compiler.

[CUCS-046-92] Imaging-consistent reconstruction/restoration and the integrating resampler.

George Wolberg and Terry Boult.

This paper introduces a new class of reconstruction/restoration algorithms that are fundamentally different from traditional approaches. We deviate from the standard practice that treats images as point samples. In this work, image values are treated as area samples generated by nonoverlapping integrators. This is consistent with the image formation process, particularly for CCD and CID cameras. We show that superior results are obtained by formulating reconstruction as a two-stage process: image restoration followed by application of the point spread function (PSF) of the imaging sensor. By coupling the PSF to the reconstruction process, we satisfy a more intuitive fidelity measure of accuracy that is based on the physical limitations of the sensor. Efficient local techniques for image restoration are derived to invert the effects of the PSF and estimate the underlying image that passed through the sensor. The reconstruction algorithms derived herein are local methods that compare favorably to cubic convolution, a well-known local technique, and they even rival global algorithms such as interpolating cubic splines. We introduce integrating resamplers as an efficient method for implementing the proposed reconstruction/restoration algorithms. Integrating resampling has been limited in the past to linear reconstruction. This paper shows how to extend it to the methods described herein. Examples are given, with applications drawn from image warping for geometric correction of images degraded by lens and chromatic aberrations.

[CUCS-047-92] Process evolution for the MARVEL environment.

Gail E. Kaiser, Israel Z. Ben-Shaul, George T. Heineman, John K. Hinsdale, and Wilfredo Marrero.

We present a schema and process evolution tool, called the Evolver, for the Marvel process-centered environment. The Evolver analyzes the differences between the new and installed process models of an existing environment, detecting each case where the notion of consistency defined by the process model has been strengthened or weakened. The Evolver then automatically updates the environment's objectbase to guarantee that the objects are consistent according to the new specifications. The Evolver can be applied while the installed process is in progress, temporarily halting normal operation while it updates the objectbase, after which development continues using the new process. We have had several months of experience using the Evolver to make repeated changes in the process that supports our own further development of Marvel, and include in this paper one small but practical example of a recent change made to a real Marvel process.

[CUCS-048-92] Pdbtool: A prototype object oriented toolkit for protein structure verification.

Calton Pu, K. P. Sheka, J. Ong, A. Chang L. Chang, and E. Alessio.

Very few research efforts have been directed towards using the object-oriented technology to solver some of the problems in the molecular biology domain. Such problems, like modeling proteins as objects and being able to manipulate them, as well as verifying its structure, are the purpose of this common research effort, which at the same time is quite unique, since it joins forces from two completely different areas: Biophysics and Computer Sciences. This report concentrates on the different steps taken to build a useful tool for the crystallographer community with an innovative development technique based on the Object Oriented paradigm.

[CUCS-049-92] Tight comparison bounds for the string prefix-matching problem.

Dany Breslauer, Livio Colussi, and Laura Toniolo.

In the *string prefix-matching* problem one is interested in finding the longest prefix of a pattern string of length m that occurs starting at each position of a text string of length n. This is a natural generalization of the string matching problem where only occurrences of the whole pattern are sought. The Knuth-Morris-Pratt string matching algorithm can be easily adapted to solve the string prefix-matching problem without making additional comparisons. In this paper we study the exact complexity of the string prefix-matching problem in the deterministic sequential comparison model. Our bounds do not account for comparisons made in a pattern preprocessing step. The following results are presented:

1. A family of linear-time string prefix-matching algorithms that make at most $\lfloor \frac{2m-1}{m} n \rfloor$ comparisons.
2. A tight lower bound of $\lfloor \frac{2m-1}{m} n \rfloor$ comparisons for any string prefix-matching algorithm.

We also consider the special case when the pattern and the text strings are the same string and all comparisons are accounted. This problem, which we call the *string prefix-function* problem, is solved in the pattern preprocessing of the Knuth-Morris-Pratt string matching algorithm using $2m-1$ comparisons. By using the lower bound for the string prefix-matching problem we are able to show:

3. A lower bound of $\lfloor 2m - \frac{m}{\lfloor\sqrt{m}\rfloor} - \lfloor\sqrt{m}\rfloor\rfloor$ comparisons for any algorithm that computes the prefix-function.

.

[CUCS-050-92] Isochronets: a high-speed network switching architecture.

Yechiam Yemini and Danilo Florissi.

This paper overviews a novel switching architecture for high-speed networks: Isochronets. Isochronets time-divide network bandwidth among routing trees. Traffic moves down a routing tree to the root during its time band. Network functions such as routing and flow control are entirely governed by band timers and require no processing of frame headers bits. Frame motions need not be delayed for switch processing, allowing Isochronets to scale over a large spectrum of transmission speeds and support all-optical implementations. The network functions as a media-access layer that can support multiple framing protocols simultaneously, handled by higher layers at the periphery. Internetworking is reduced to a simple media-layer bridging. Isochronets provide flexible quality of service control and multicasting through allocation of bands to routing trees. They can be tuned to span a spectrum of performance behaviors outperforming both circuit or packet switching.

[CUCS-051-92] On hard real-time management information.

Shyhtsun F. Wu and Gail E. Kaiser.

Maintaining the quality of services for time-critical broadband applications on a high-speed network is an important subject for network management. Since the managed services are time-sensitive, the management applications managing these services must be *real-time*. This further implies that the real-time management applications must first have real-time management information to correctly take management actions. In this paper, we investigate how to build a monitoring system to support *hard* real-time management information in the domain of network management. The key issue to design such a hard real-time system is to *bound* the amount of *resources* that would be needed at the time of monitoring. Three types of resources are under our consideration: communication, computation, and information. This resource bound property motivates us to extend the existing management information protocols (e.g., SNMP or CMIP) to achieve hard real-time monitoring. We describe our experience in building a *soft* real-time network monitoring system called MeldNet, which uses the *shadow object* feature provided by the MeldC distributed concurrent object-oriented programming system. Then, we introduce the *DONUT* real-time scheduling model, which is designed to solve real-time scheduling problems in the domain of hard real-time monitoring. Finally, we discuss the negotiation process between the management applications and the hard real-time monitoring system.

[CUCS-052-92] Saving comparisons in the crochemore-perrin string matching algorithm.

Dany Breslauer.

Crochemore and Perrin discovered an elegant linear-time constant-space string matching algorithm that makes at most $2n-m$ symbol comparison. This paper shows how to modify their algorithm to use fewer comparisons. Given any fixed $\epsilon > 0$, the modified algorithm takes linear time, uses constant space and makes at most $n + \lfloor\frac{1+\epsilon}{2}(n-m)\rfloor$ comparisons. If $O(\log m)$ space is available, then the algorithm makes at most $n + \lfloor\frac{1}{2}(n-m)\rfloor$ comparisons. The pattern preprocessing step also takes linear time and uses constant space. These are the first string matching algorithms that make fewer than $2n-m$ comparisons and use sub-linear space.

[CUCS-053-92] Testing string superprimitivity in parallel.

Dany Breslauer.

A string w *covers* another string z if every symbol of z is within some occurrence of w in z. A string is called *superprimitive* if it is covered only by itself, and *quasiperiodic* if it is covered by some shorter string. We present an $O(\log\log n)$ time $\frac{n\log n}{\log\log n}$-processor CRCW-PRAM algorithm that tests if a string is superprimitive.

[CUCS-054-92] Efficient comparison based string matching.

Dany Breslauer and Zvi Galil.

We study the exact number of symbol comparisons that are required to solve the string matching problem and present a family of efficient algorithms. Unlike previous string matching algorithms, the algorithms in this family do not "forget" results of comparisons, what makes their analysis much simpler. In particular, we give a linear-time algorithm that finds all occurrences of a pattern of length m in a text of length n in $n + \lceil\frac{4\log m+2}{m}(n-m)\rceil$ comparisons. The pattern preprocessing takes linear time and makes at most $2m$ comparisons. This algorithm establishes that, in general, searching for a long pattern is easier than searching for a short one. We also show that any algorithm in the family of the algorithms presented must make at least $n + \lfloor\log m\rfloor\lfloor\frac{n-m}{m}\rfloor$ symbol comparisons, for $m = 2^k - 1$ and any integer $k \geq 1$.

[CUCS-055-92] Reflectance ratio: A photometric invariant for object recognition.

Shree K. Nayar and Ruud M. Bolle.

This paper presents a photometric invariant called reflectance ratio that can be computed from a single brightness image of a scene. Neighboring points on a smoothly curved surface have similar surface orientations and illumination conditions. Hence, their brightness values can be used to compute the ratio of their reflectance coefficients. Based on this observation, we develop an algorithm that estimates a reflectance ratio for each region in an image with respect to its background. The algorithm is computationally efficient as it computes ratios for all image regions in just two raster scans. The region reflectance ratio represents a physical property of a region that is invariant to the illumination conditions. Several experimental results are included to demonstrate the accuracy and robustness of reflectance ratios. We use the reflectance ratio invariant to recognize objects from a single brightness image of a scene. Object models are automatically acquired and represented using a hash table. Recognition and pose estimation algorithms are presented that use the reflectance ratios of scene regions as well as their geometric properties to index the hash table. The result is a hypothesis for the existence of an object in the image. This hypothesis is verified using the reflectance ratios and locations of other regions in the scene. This approach to the recognition problem is very effective for objects with printed characters and pictures. We conclude with experimental results on recognition that demonstrate the power of using photometric and geometric properties of objects, simultaneously.

[CUCS-056-92] Illumination planning in parametric eigenspace.

Hiroshi Murase and Shree K. Nayar.

We present a new approach to the problem of illumination planning. Given a set of objects, the goal is to determine the illumination direction for which the objects are most distinguishable in appearance from each other. Correlation is used as a measure of the similarity in the appearance of objects. For each object, a large number of images are automatically obtained by varying pose and illumination direction. The images set of all objects, together, constitute the planning set. The planning set is compressed using the Karhunen-Loeve transform to obtain a low-dimensional eigenspace. For each illumination direction, objects are represented as hypercurves in the eigenspace. The minimum distance between the hypercurves of two objects, represents the similarity between the objects in the correlation sense. The optimal source direction is one that maximizes the shortest distance between object hypercurves in eigenspace. Several experiments have been conducted using real objects. Results produced by the illumination planner have been used to improve the performance of an object recognition system. The illumination planning method presented here has the following advantages over previous methods. (a) The proposed method does not rely on geometric (CAD) models of the objects of interest. It uses only brightness images to accomplish the planning task. (b) No assumptions are made with respect to the reflectance properties of the objects. (c) The optimal source direction produced by the illumination planner is pose invariant; it is optimal over all object poses.

[CUCS-057-92] Diffuse reflectance from rough surfaces.

Michael Oren and Shree K. Nayar.

The Lambertian assumption is one of the most widely used assumptions in machine vision. It is used explicitly by vision techniques such as shape from shading and photometric stereo, and implicitly by techniques such as binocular stereo and motion detection. We show that rough surfaces, even when locally Lambertian, are non-Lambertian in reflectance. While the radiance of a Lambertian surface is independent of viewing direction, the radiance of a rough diffuse surface increases as the viewing direction approaches the source direction. This preferred scattering around the source direction results primarily from the tilt of points on the rough surface with respect to the global orientation of the surface. The reflectance model presented here accounts for complex geometric and radiometric phenomena such as masking, shadowing, and interreflections between points on the rough surface. This new model may be viewed as a generalization of the Lambertian model. We have conducted several experiments on samples of rough diffuse surfaces, such as, plaster and sand. All of these samples demonstrate significant deviation from Lambertian behavior. The reflectance measurements obtained are in strong agreement with the reflectance predicted by the model. We conclude with a brief discussion on the implications of these results on machine vision. We present a new approach to the problem of illumination planning. Given a set of objects, the goal is to determine the illumination direction for which the objects are most distinguishable in appearance from each other. Correlation is used as a measure of the similarity in the appearance of objects. For each object, a large number of images are automatically obtained by varying pose and illumination direction. The images set of all objects, together, constitute the planning set. The planning set is compressed using the Karhunen-Loeve transform to obtain a low-dimensional eigenspace. For each illumination direction, objects are represented as hypercurves in the eigenspace. The minimum distance between the hypercurves of two objects, represents the similarity between the objects in the correlation sense. The optimal source direction is one that maximizes the shortest distance between object hypercurves in eigenspace. Several experiments have been conducted using real objects. Results produced by the illumination planner have been used to improve the performance of an object recognition system. The illumination planning method presented here has the following advantages over previous methods. (a) The proposed method does not rely on geometric (CAD) models of the objects of interest. It uses only brightness images to accomplish the planning task. (b) No assumptions are made with respect to the reflectance properties of the objects. (c) The optimal source direction produced by the illumination planner is pose invariant; it is optimal over all object poses.

[CUCS-058-92] Removal of specularities using color and polarization.

Shree K. Nayar, Xi-Sheng Fang, and Terrance Boult.